OptQC: An optimised parallel quantum compiler

T. Loke, J. B. Wang [1] and Y.H. Chen

School of Physics, The University of Western Australia, 6009 Perth, Australia

arXiv:1408.1612v1 [quant-ph] 7 Aug 2014

Abstract

The software package *Qcompiler* [1] provides a general quantum compilation framework, which maps any given unitary operation into a quantum circuit consisting of a sequential set of elementary quantum gates. In this paper, we present an extended software *OptQC*, which finds permutation matrices P and Q for a given unitary matrix U such that the number of gates in the quantum circuit of $U = Q^T P^T U' PQ$ is significantly reduced, where U' is equivalent to U up to a permutation and the quantum circuit implementation of each matrix component is considered separately. We extend further this software package to make use of high-performance computers with a multiprocessor architecture using MPI. We demonstrate its effectiveness in reducing the total number of quantum gates required for various unitary operators.

Program summary *Program title*: *OptQC*

Catalogue identifier:

Program summary: URL: http://cpc.cs.qub.ac.uk/summaries/

Program obtainable from: CPC Program Library, Queens University, Belfast, N. Ireland

Licensing provisions: Standard CPC licence, http://cpc.cs.qub.ac.uk/licence/licence.html

Distribution format: tar.gz

Programming language: Fortran, MPI

Computer: Any computer with Fortran compiler and MPI library.

Operating system: Linux

Classification:

Nature of problem: It aims to minimise the number of quantum gates required to implement a given unitary operation.

[1]corresponding author: jingbo.wang@uwa.edu.au

Solution method: It utilises a threshold-based acceptance strategy for simulated annealing to select permutation matrices P and Q for a given unitary matrix U such that the number of gates in the quantum circuit of $U = Q^T P^T U' P Q$ is minimised, where U' is equivalent to U up to a permutation. The decomposition of a unitary operator is performed by recursively applying the cosine-sine decomposition.

Running time: Running time increases with the size of the unitary matrix, as well as the prescribed maximum number of iterations for qubit permutation selection and the subsequent simulated annealing algorithm. Running time estimates are provided for each example in Section 4. All simulation results presented in this paper are obtained from running the program on the Fornax supercomputer managed by iVEC@UWA with Intel Xeon X5650 CPUs.

---

1. Introduction

Quantum computation aims to solve problems that are classically intractable by harnessing intricate quantum correlations between densely encoded states in quantum systems [2]. A well-known example is Shor's algorithm for the factorisation of numbers [3, 4]. Quantum algorithms are designed to be implemented on quantum computers by means of a quantum circuit, which consists of qubits and quantum gates. It is therefore of vital importance to be able to obtain a quantum circuit representation for any given quantum algorithm (which is always described by a unitary matrix) in terms of an elementary set of quantum gates - this role is that of a quantum compiler.

Barenco *et al* and Deutsch *et al* [5, 6] proved that any arbitrarily complex unitary operation can be implemented by a quantum circuit involving only one- or two-qubit elementary quantum logic gates. Earlier studies applied the standard triangularisation or QR-factorisation scheme with Givens rotations and Gray codes to map a quantum algorithm to a series of elementary gate operations [2, 5–7]. Several research groups examined a more efficient and versatile scheme based on the cosine-sine decomposition was proposed and utilized [8–12]. De Vos *et al* [13, 14] looked into another decomposition scheme, namely the Birkhoff decomposition, which was found to provide simpler quantum circuits for certain types of unitary matrices than the cosine-sine decomposition. However, the Birkhoff decomposition does not work for general unitary matrices.

More recently, Chen and Wang [1] developed a general quantum compiler package written in Fortran, entitled the *Qcompiler*, which is based on the cosine-sine decomposition scheme and works for arbitrary unitary matrices. The number of gates required to implement a general 2^n-by-2^n unitary matrix using the CSD method scales as $O(4^n)$ [9, 12]. In other words, the number of gates scales exponentially with the number of qubits. Thus, in any practical application of the CSD method to decomposing matrices, it is of considerable interest to reduce the number of gates required as much as possible.

In this work, we adopt the CSD method due to the reasons outlined above, and we split the unitary matrix U into an equivalent sequence of unitaries with the aim of reducing the number of gates required to implement the entire sequence of unitaries. In general, this

means writing U as a sequence of s unitaries, i.e. $U = U_s U_{s-1} \dots U_1$. At first glance, this seems counterintuitive, since if we were to apply the CSD to each unitary, this would increase the scaling of the number of gates required to $\mathrm{O}(4^n s)$, which is undesirable. However, we note that (1) certain U_i can be decomposed more efficiently than CSD such as qubit permutation matrices; and (2) some matrices requires only a few gates when separately decomposed using the CSD method.

This paper is organised as follows. Section 2 describes in detail our approach for reducing the number of gates required to implement any given unitary matrix U. Section 3 details our developed program, called *OptQC*, that uses the methods described in section 2 to reduce the number of gates required to implement any given unitary matrix. Some sample results using the program are given in section 4, and then we discuss our conclusions and possible future work in section 5.

2. Our approach

Suppose we are given an m-by-m (where $m = 2^n$) unitary matrix U. As mentioned above, we are interested in splitting U into a sequence of unitaries with the aim of reducing the total number of gates required to implement the entire sequence. One means of splitting U into an equivalent sequence of unitaries is by using permutation matrices. A permutation matrix is a square binary matrix that contains, in each row and column, precisely a single 1 with 0s everywhere else. For any permutation matrix P, its corresponding inverse is $P^{-1} = P^T$. For convenience, we also define an equivalent representation of permutations using lists - a permutation list p (lowercase) is equivalent to the permutation matrix P (uppercase) by the relation:

$$(P)_{i,j} = \delta_{p[\![i]\!],j}, \tag{1}$$

where δ is the Kronecker delta function, and $p[\![i]\!]$ denotes the ith list element of p. For example, the permutation list $p = \{2, 1, 4, 3\}$ corresponds to the 4-by-4 permutation matrix:

$$P = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}.$$

Now define $\mathrm{CSD}(M)$ to be the number of gates required to implement the unitary matrix M according to the CSD method. If we were to write U as $U = P^T U' P$ (where U' is equivalent to U up to a permutation), then we find that $\mathrm{CSD}(U) \neq \mathrm{CSD}(U') + \mathrm{CSD}(P) + \mathrm{CSD}(P^T)$ in general (note also that $\mathrm{CSD}(P) \neq \mathrm{CSD}(P^T)$). The general aim is thus to find a P that minimises the total cost function $\mathrm{CSD}(U') + \mathrm{CSD}(P) + \mathrm{CSD}(P^T)$, with the obvious restriction that it has to be less than $\mathrm{CSD}(U)$.

In our approach, we write U as $U = Q^T P^T U' P Q$, where P and Q are both permutation matrices, and U' is equivalent to U up to a permutation. In general, P is allowed to be any permutation matrix ($m! = (2^n)!$ permutations possible), but Q is restricted to a class of permutation matrices that correspond to qubit permutations (only $n!$ permutations possible).

The advantage of this approach is that qubit permutations can be easily implemented using a sequence of swap gates - so for a system with n qubits, it requires at most $n-1$ swap gates to implement any qubit permutation. This also enables the program (in the parallel version) to start different threads at different points in the search space of $m!$ permutations by using different qubit permutations.

Let $s_{num}(Q)$ be the number of swap gates required to implement a qubit permutation matrix Q. Note that $s_{num}(Q) = s_{num}(Q^T)$, since the reverse qubit permutation would just be the same swap gates applied in reverse order. The total cost function c_{num} of implementing a given unitary $U = Q^T P^T U' P Q$ is then:

$$c_{num}(U) = \mathrm{CSD}(U') + \mathrm{CSD}(P) + \mathrm{CSD}(P^T) + 2s_{num}(Q). \tag{2}$$

To make the dependencies in this function clear, we write this as:

$$c_{num}(U, P, Q) = \mathrm{CSD}(PQUQ^T P^T) + \mathrm{CSD}(P) + \mathrm{CSD}(P^T) + 2s_{num}(Q), \tag{3}$$

which is the function that we aim to minimise with respect to P and Q.

3. Program outline

We have developed a Fortran program, called *OptQC*, which reads in a unitary matrix U, minimises the total cost function $c_{num}(U, P, Q)$, and outputs a quantum circuit that implements U. A significant portion of this program is based on the CSD code provided by the LAPACK library [15]) and the recursive procedure implemented in *Qcompiler*, developed by Chen and Wang [1]. As with *Qcompiler*, the new *OptQC* program has two different branches, one treating strictly real unitary (i.e. orthogonal) matrices, and another treating arbitrary complex unitary matrices, with the former generally providing a circuit that is half in size of the latter [1].

Note that the CSD procedure requires the round up of the matrix dimension to the closest power of two, i.e. the dimension used is

$$m' = 2^{\lceil log_2 m \rceil}. \tag{4}$$

The expanded unitary operator $\bar{U}$ is an m'-by-m' matrix, where

$$\left(\bar{U}\right)_{i,j} = \begin{cases} (U)_{i,j} & : i \leq m, j \leq m \\ \delta_{i,j} & : \text{otherwise} \end{cases} \tag{5}$$

which we will subsequently treat as the unitary U to be optimised via permutations.

In the following subsections we describe the key procedures in *OptQC*, depicted in Figure 1, which serve to progressively reduce the total cost function $c_{num}(U, P, Q)$. We first detail the serial version of the program, followed by an extension to a parallel architecture using MPI.

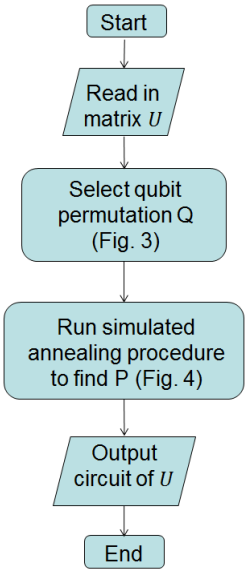
Start
Read in matrix U
Select qubit permutation Q (Fig. 3)
Run simulated annealing procedure to find P (Fig. 4)
Output circuit of U
End

Figure 1: Flowchart overview of the serial version of *OptQC*.

3.1. Selection of qubit permutation

Qubit permutations are a class of permutations that are expressible in terms of a reordering of qubits, which can be efficiently implemented using swap gates that serve to interchange qubits. Recalling that U is of dimensions m-by-m (where $m = 2^n$), this implies that there are only $n!$ qubit permutations possible for a given U. A qubit permutation can be expressed as a list q (lowercase) of length n, or as a permutation matrix Q (uppercase) of dimensions m-by-m. A qubit permutation of length n requires at most $n-1$ swap gates.

The selection of the qubit permutation matrix Q is done by varying Q and computing the corresponding change in the cost function c_{num}, while holding P constant as the identity matrix I. An example implementation of the $n = 3$ qubit permutation $q = \{3, 1, 2\}$ is shown in Figure 2. By considering how the basis states are mapped to each other by q, a regular permutation list $\bar{q}$ of length m can be readily constructed from q, and then we use the relation between permutation lists and permutation matrices (see Eq. (1)) to obtain Q from $\bar{q}$.

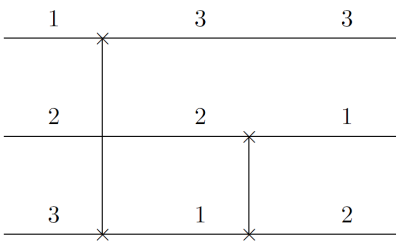

Figure 2: Example implementation of qubit permutation $q = \{3, 1, 2\}$ using swap gates.

We start the program with an identity qubit permutation q, i.e. $q[[i]] = i$ (corresponding

to $Q = I$), and compute the corresponding cost of implementation $c_{num}(U, I, Q)$. Then, for some prescribed number of iterations j_{max}, we generate a random qubit permutation q' each time and compute the new cost as $c_{num}(U, I, Q')$. If the new cost is lower than the initial cost (recorded by $c_{num}(U, I, Q)$), the current qubit permutation q is replaced by q'. Figure 3 shows a flowchart overview of the qubit selection procedure. After this procedure, we have an optimised qubit permutation matrix Q, which will remain unchanged while we find the unrestricted permutation matrix P in the next section through a simulated annealing process.

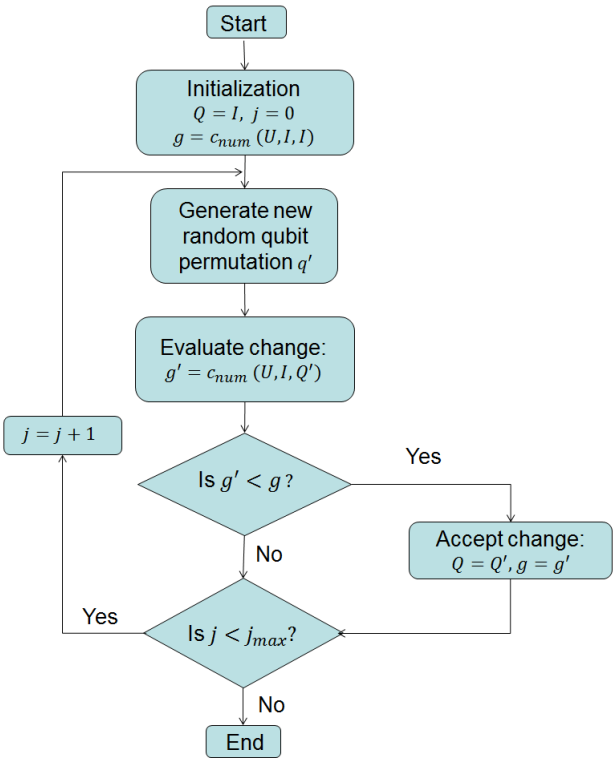

Figure 3: Flowchart overview of the qubit selection procedure.

3.2. Simulated annealing

Here, we aim to find an optimal permutation p' such that $c_{num}(U, P', Q) < c_{num}(U, P, Q)$ in the discrete search space of all $m!$ permutations. Given the massive size of the search space, use of a heuristic optimisation method is practically necessary. Simulated annealing is one such method for finding a minimum in a discrete search space. In the OptQC program, we adopt a threshold acceptance based simulated annealing algorithm. There are three key components to the algorithm:

1. Cost function: the function to be minimised, i.e. $c_{num}(U, P, Q)$.

2. Neighbourhood operator: the procedure that alters the current solution slightly by altering the current permutation p to a slightly different permutation p'. Our neighbourhood operator acts to interchange any two *random* positions in p to form p'.

3. Threshold value: any 'bad' trades (increase in cost function) that are below some threshold value β are accepted, otherwise they are rejected. We define the threshold value as $\beta(P,Q) = \min\left(\lceil \alpha c_{num}(U,P,Q) \rceil, \lceil \alpha c_{num}(U,I,Q) \rceil\right)$, where $0 \leq \alpha < 1$. As such, the threshold value is taken to be the proportion α of the current number of gates (with a fixed maximum value of the proportion α of the initial number of gates to ensure that $\beta(P,Q)$ cannot grow arbitrarily large).

We start with p as the identity permutation. By iterating the neighbourhood operator and evaluating the subsequent change in the number of gates, we accept the change in the permutation if it reduces the number of gates, or if the increase in the number of gates is below the threshold β. After some prescribed number of iterations i_{max}, we terminate the simulated annealing procedure, returning the permutation p_{min} that provides the minimum number of gates. Figure 4 shows a flowchart overview of the simulated annealing procedure. Note that p_{min} is not necessarily the permutation p at the end of i_{max} iterations - rather, we keep track of p_{min} separately during the procedure.

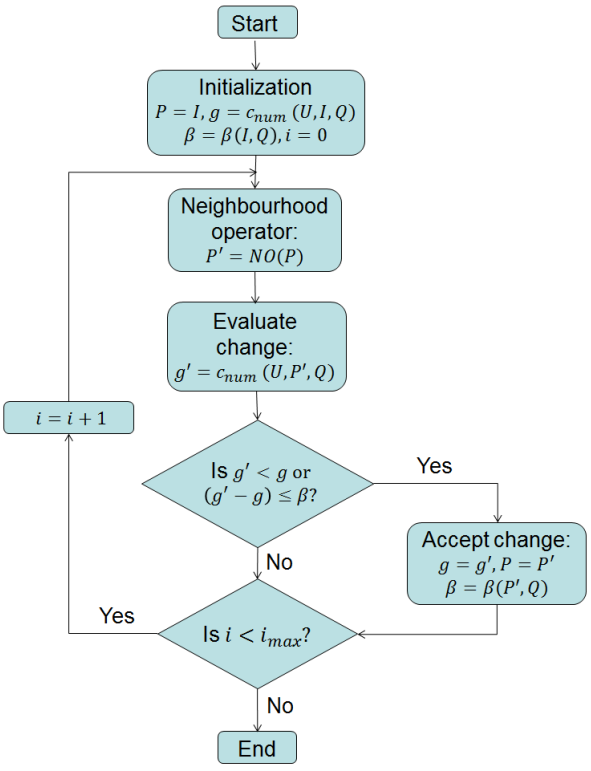

Figure 4: Flowchart overview of the simulated annealing procedure.

3.3. Gate reduction procedure

Here, we focus on reducing the number of gates in some prescribed quantum circuit by combining 'similar' gates. In a quantum circuit, we can combine CUGs (controlled unitary gates) that apply the same unitary operation U_{op} to the same qubit, with all but one of the conditionals of the CUGs being the same. This reduction process is carried out after every application of the CSD method to a matrix - in particular, it is applied three times (to $PQUQ^TP^T$, P and P^T respectively) when computing $c_{num}(U, P, Q)$ (see Eq. (3)). While it does impose a significant computational overhead, it gives a better reflection of the true cost function, since the reduced circuit is the circuit that one would use for implementation. Figure 5 shows an example result of applying the reduction procedure to a quantum circuit.

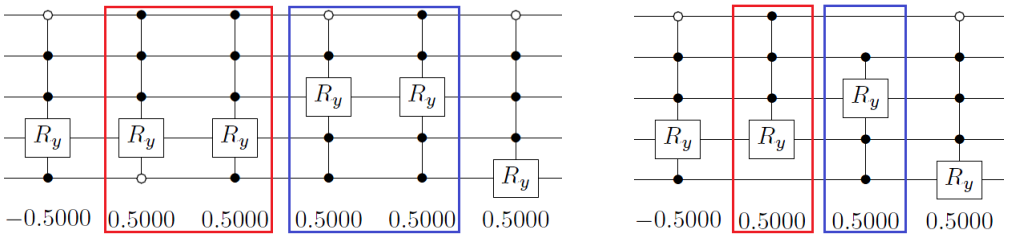

(a) Circuit before reduction procedure

(b) Circuit after reduction procedure

Figure 5: Example of applying the reduction procedure to a quantum circuit.

3.4. MPI parallelisation

The program described above can be readily extended to a parallel architecture using MPI. Since the neighbourhood operator in the simulated annealing procedure acts to interchange any two random positions, it follows that if the random number generator is seeded differently, then a different set of positions would be interchanged, i.e. a different search through the space of permutations would be conducted. Similarly, the qubit permutation that is generated would also change when seeded differently, which enables the program to start threads at multiple locations in the search space of $m!$ permutations, so that the search procedure explores as much of the permutation space as possible. We do, however, restrict the root thread (thread index 0) of the program to use the identity qubit permutation for comparison purposes. Hence, using MPI, we can spawn a team of threads that simultaneously searches through the space of permutations independently and differently (by seeding the random number generator of each thread differently), and then collate the results to pick out the thread with the most optimal permutation, that is, it has the lowest $c_{num}(U, P, Q)$ value.

4. Results

We now apply the software program *OptQC* to various unitary operations to obtain corresponding optimised quantum circuits. All the results shown here are obtained using

parameters $i_{max} = 40000$, $j_{max} = 1000$ and $\alpha = 0.01$ (we choose this α value because it provides, on average, the best results for the unitary operators being considered in this paper. Using these parameters, we run *OptQC* on the supercomputer Fornax with Intel Xeon X5650 CPUs, managed by iVEC@UWA, using 8 nodes with 12 cores on each (i.e. 96 threads).

4.1. Real unitary matrix

A random real unitary (i.e. orthogonal) matrix is given below:

$$U = \begin{pmatrix} 0.0438 & 0 & 0 & 0 & 0.9990 & 0 & 0 & 0 \\ 0.1297 & 0.8689 & -0.2956 & 0 & -0.0057 & 0.1538 & -0.3423 & 0 \\ -0.2923 & 0 & 0.6661 & 0 & 0.0128 & 0 & -0.6861 & 0 \\ -0.0061 & -0.0412 & 0.0140 & 0.7058 & 0.0003 & 0.3008 & 0.0162 & -0.6397 \\ 0.9147 & 0 & 0.4021 & 0 & -0.0401 & 0 & 0 & 0 \\ 0.0185 & 0.1242 & -0.0422 & 0.3961 & -0.0008 & -0.9073 & -0.0489 & 0 \\ 0.2424 & -0.4762 & -0.5524 & 0 & -0.0106 & 0 & -0.6397 & 0 \\ 0.0051 & 0.0343 & -0.0117 & -0.5874 & -0.0002 & -0.2503 & -0.0135 & -0.7686 \end{pmatrix}$$

Note that this matrix is not completely filled, otherwise no reduction via permutations would generally be possible. By using *OptQC*, the reduction process gives the following results for the thread which achieves the optimal solution:

- No optimisation: $c_{num}(U, I, I) = 29$ gates
- After selection of an optimised qubit permutation q: $c_{num}(U, I, Q) = 1 + 26 + 1 = 28$ gates
- After simulated annealing process for the permutation p: $c_{num}(U, P, Q) = 1 + 2 + 16 + 2 + 1 = 22$ gates

Hence, we achieve a reduction of $\sim 25\%$ from the original number of gates. Figure 6 shows a comparison between the original and optimised circuit for U. Runtime for this calculation is ~ 14.5 seconds.

4.2. Quantum walk operators

One important class of unitary operators are quantum walk operators - in particular, discrete-time quantum walk (DTQW) step operators [16–18]. For a given undirected graph $G(V, E)$, defined by a vertex set V and edge set E, we can define the DTQW step operator $U = SC$, where S and C are the shifting and coin operators respectively. The shifting operator acts to swap coin states that are connected by an edge, and the coin operator acts to mix the coin states at each individual vertex.

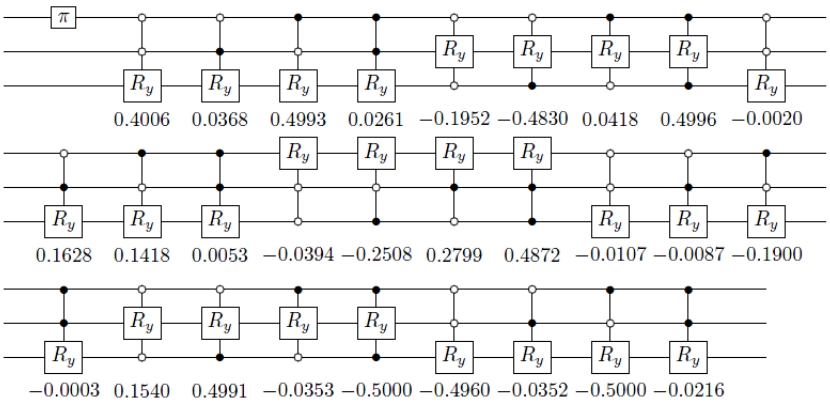

(a) Original circuit - 29 gates

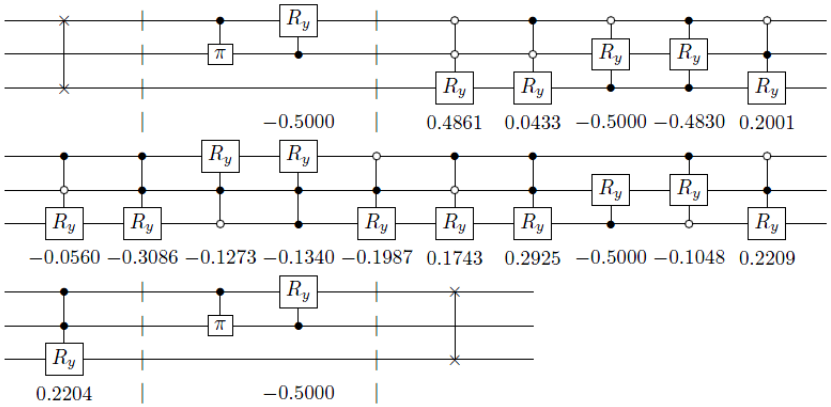

(b) Optimised circuit - 22 gates

Figure 6: Result of quantum circuit optimisation as performed by *OptQC* on a random real unitary matrix. In (b), the dashed vertical lines separate the circuit for each matrix - from left to right, this corresponds to Q, P, U', P^T and Q^T respectively.

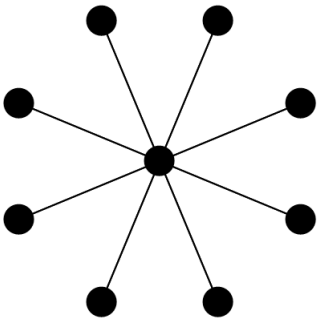

Figure 7: The 8-star graph

4.2.1. 8-star graph

The 8-star graph (shown in Figure 7) is a graph with 1 centre vertex connected to 8 leaf vertices by undirected edges. Using the Grover coin operator, the resulting quantum walk operator on this graph corresponds to a 16-by-16 real unitary matrix, as given below:

$$U = \begin{pmatrix}
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -0.75 & 0.25 & 0.25 & 0.25 & 0.25 & 0.25 & 0.25 & 0.25 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.25 & -0.75 & 0.25 & 0.25 & 0.25 & 0.25 & 0.25 & 0.25 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.25 & 0.25 & -0.75 & 0.25 & 0.25 & 0.25 & 0.25 & 0.25 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.25 & 0.25 & 0.25 & -0.75 & 0.25 & 0.25 & 0.25 & 0.25 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.25 & 0.25 & 0.25 & 0.25 & -0.75 & 0.25 & 0.25 & 0.25 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.25 & 0.25 & 0.25 & 0.25 & 0.25 & -0.75 & 0.25 & 0.25 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.25 & 0.25 & 0.25 & 0.25 & 0.25 & 0.25 & -0.75 & 0.25 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.25 & 0.25 & 0.25 & 0.25 & 0.25 & 0.25 & 0.25 & -0.75 \\
1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0
\end{pmatrix}$$

By using *OptQC*, the reduction process gives the following results for the thread which achieves the optimal solution:

- No optimisation: $c_{num}(U, I, I) = 34$ gates
- After selection of an optimised qubit permutation: $c_{num}(U, I, Q) = 27$ gates
- After simulated annealing process to select a permutation p: $c_{num}(U, P, Q) = 0 + 2 + 19 + 2 + 0 = 23$ gates

Hence, we achieve a reduction of $\sim 32\%$ from the original number of gates. Figure 6 shows the optimized circuit obtained for U. Runtime for this calculation is ~ 47 seconds.

4.2.2. 3rd generation 3-Cayley tree

The 3rd generation 3-Cayley tree (abbreviated as the 3CT3 graph) is a tree of 3 levels in which all interior nodes have degree 3, as shown in Figure 9(a). The corresponding quantum walk operator using the Grover coin operator is shown in Figure 9(b) - the quantum walk operator U is a 42-by-42 real unitary matrix (which is fairly sparse), which, for the purposes of the decomposition, is expanded to a 64-by-64 unitary matrix as per Eq. (5).

By using *OptQC*, the reduction process gives the following results for the thread which achieves the optimal solution:

- No optimisation: $c_{num}(U, I, I) = 996$ gates

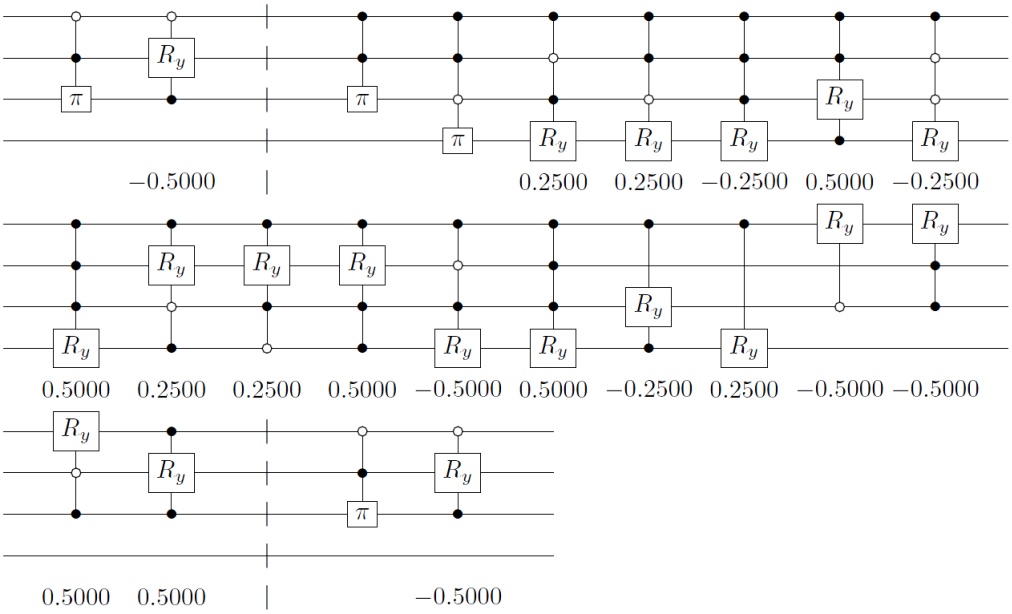

Figure 8: Optimised circuit (with 23 gates) for the quantum walk operator of the 8-star graph.

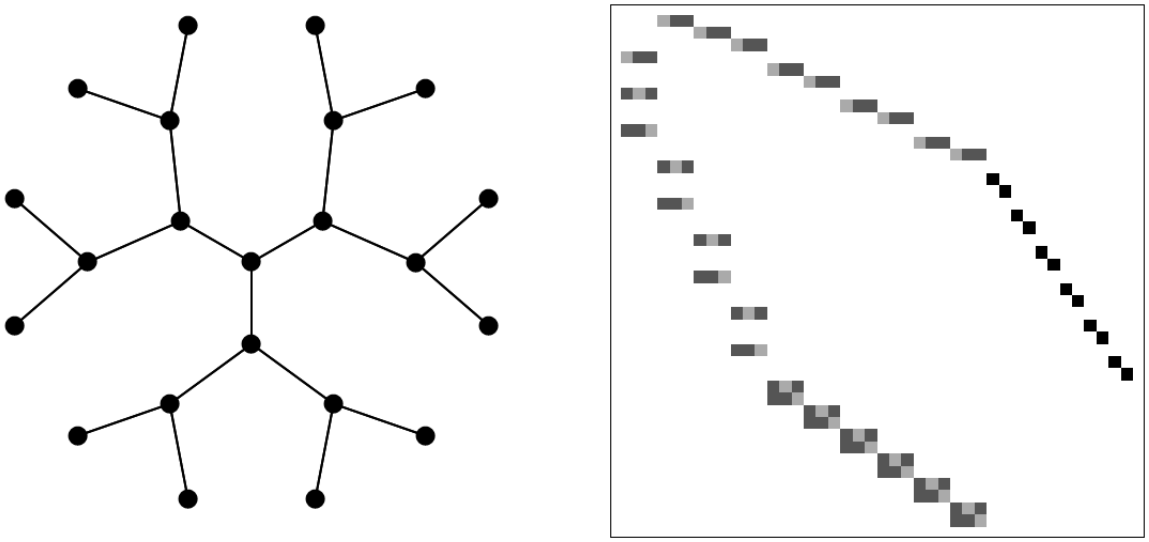

(a) 3CT3 graph

(b) Quantum walk operator for the 3CT3 graph

Figure 9: The 3CT3 graph and its corresponding quantum walk operator using the Grover coin operator. The colours/shades in (b) denote the matrix entries for $-1/3$ (light grey), $2/3$ (dark grey) and 1 (black) - all other matrix entries are 0 (white).

- After selection of an optimised qubit permutation: $c_{num}(U, I, Q) = 3 + 345 + 3 = 351$ gates
- After simulated annealing process to select a permutation p: $c_{num}(U, P, Q) = 3 + 33 + 231 + 30 + 3 = 300$ gates

Hence, we achieve a reduction of $\sim 70\%$ from the original number of gates. Runtime for this calculation is $\sim$ 12 minutes. Figure 10 shows the time-series for $c_{num}(U, P, Q)$ during both the qubit permutation selection phase and the simulated annealing process (separated by a dotted line) to achieve the above result.

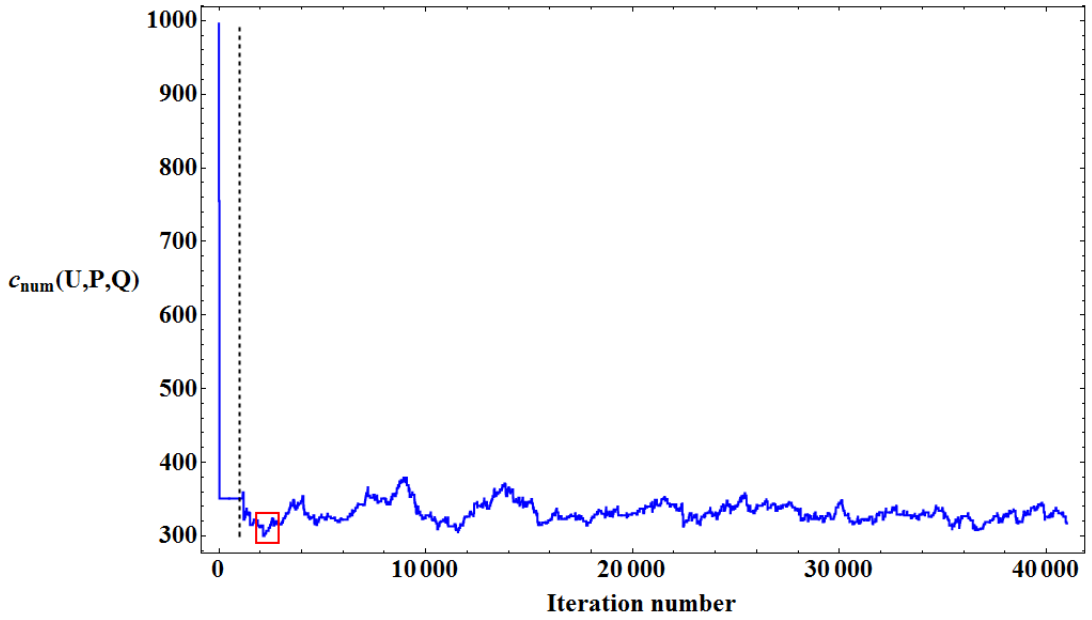

Figure 10: Time-series of $c_{num}(U, P, Q)$ during the simulated annealing process for the thread which obtains the optimal solution. The original number of gates required by the CSD method to implement U is 996; selection of a qubit permutation reduces this cost to 351 gates, which is used as the starting point for the simulated annealing process. The red box indicates the region where the optimal solution of 300 gates is achieved. The iterations before the dotted line indicate the qubit permutation selection phase, and the subsequent iterations shows the simulated annealing process.

4.3. Quantum Fourier transform

Quantum Fourier transform is the quantum counterpart of the discrete Fourier transform in classical computing. It is an essential ingredient in several well-known quantum algorithms, such as Shor's factorization algorithm [3] and the quantum phase estimation algorithm [19]. The matrix representation of the quantum Fourier transform on n dimensions is given by:

$$(\mathrm{QFT})_{jk} = \frac{1}{\sqrt{n}}\omega^{jk}, \quad \text{where } \omega = \exp(2\pi i/n). \tag{6}$$

An efficient quantum circuit implementation of the quantum Fourier transform is given in [2], which scales logarithmically as $O(\log(n)^2)$. Such a circuit implementation for $n = 2^6 = 64$ is shown in Figure 11.

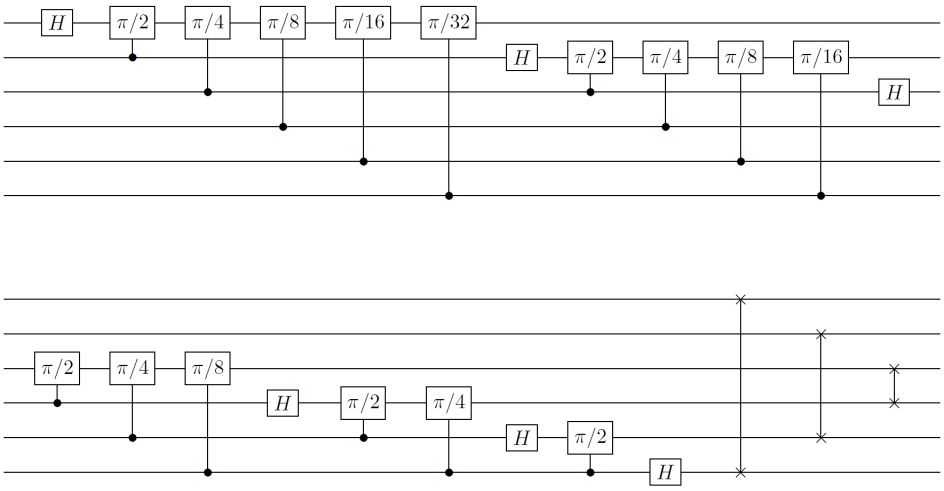

Figure 11: Circuit implementation of quantum Fourier transform for $n = 64$.

Now, let us apply *OptQC* to the corresponding 64-by-64 complex unitary operator, given by Eq. (6). With $\alpha = 0.002$, the reduction process gives the following results for the thread achieving an optimal solution:

- No optimisation: $c_{num}(U, I, I) = 4095$ gates
- After selection of an optimised qubit permutation: $c_{num}(U, I, Q) = 5+3577+5 = 3587$ gates
- After simulated annealing process to select a permutation p: $c_{num}(U, P, Q) = 5 + 69 + 3359 + 70 + 5 = 3508$ gates

Hence, we achieve a reduction of $\sim 14\%$ from the original number of gates. Runtime of this calculation is ~ 20 minutes. Figure 12 shows the time-series for $c_{num}(U, P, Q)$ during both the qubit permutation selection phase and the simulated annealing process (separated by a dotted line) to achieve the above result. Clearly, this result is by far inferior to the quantum circuit of only 24 gates shown in Figure 11. Similarly, the *OptQC* package would not be able to provide quantum circuits as efficient as those presented in [18, 20] for the implementation of quantum walks on highly symmetric graphs. This is to be expected, since the CS decomposition is a general technique that decomposes a given unitary into a fixed circuit structure using many conditional gates, with an upper bound of $\mathrm{O}(4^n)$. This algorithm is performed without foreknowledge or explicitly exploiting the structure of the unitary, which would clearly be crucial in achieving the lowest possible number of gates for a given unitary, as exemplified by the above examples. Instead, the *OptQC* package is designed to work for any arbitrary unitary operator for which we do not already have an efficient quantum circuit implementation of, for example, quantum walk operators on

arbitrarily complex graphs. In such cases, we have demonstrated that the *OptQC* package provides optimised quantum circuits that are far more efficient than the original *Qcompiler*.

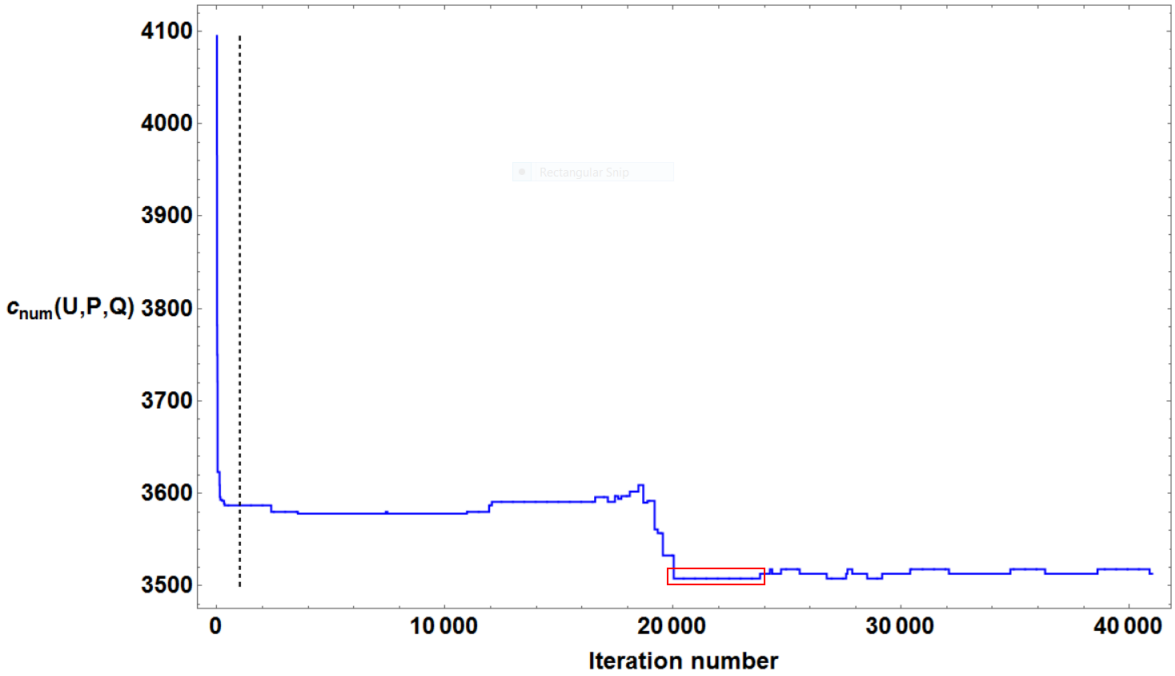

Figure 12: Time-series of $c_{num}(U, P, Q)$ during the simulated annealing process for the thread which obtains the optimal solution. The original number of gates required by the CSD method to implement U is 4095; selection of a qubit permutation reduces this cost to 3587 gates, which is used as the starting point for the simulated annealing process. The red box indicates the region where the optimal solution of 3508 gates is achieved. The iterations before the dotted line indicate the qubit permutation selection phase, and the subsequent iterations shows the simulated annealing process.

5. Conclusion and future work

We have developed an optimised quantum compiler, named as *OptQC*, that runs on a parallel architecture to minimise the number of gates in the resulting quantum circuit of a unitary matrix U. This is achieved by finding permutation matrices Q and P such that $U = Q^T P^T U' P Q$ requires less total number of gates to be implemented, where the implementation for each matrix is considered separately. Decompositions of unitary matrices is done using the CSD subroutines provided in the LAPACK library [15] and adapted from *Qcompiler* [1]. *OptQC* utilises an optimal selection of qubit permutations Q, a simulated annealing procedure to find P, and a combination of similar gates in order to reduce the total number of gates required as much as possible. We find that for many different types of unitary operators, *OptQC* is able to reduce the number of gates required by a significant amount, but its efficacy does vary depending on the unitary matrix given. In particular, this optimisation procedure works well for sparse unitary matrices.

For future work, we hope to look at characterising the optimal solutions reached to see if the matrix U' (and the associated permutation P) have some common preferential structure that leads to a reduced cost of implementation using the CSD method. Such information could be used to implement a guided search for the optimal solution, rather than using random adjustments of the permutation matrix. We also want to characterise 'bad' permutations (that is, permutations with a large cost) and avoid them in the search procedure, perhaps by eliminating the conjugacy class of 'bad' permutations from the search space.

6. Acknowledgements

Our work was supported through the use of advanced computing resources located at iVEC@UWA, as well as funding for a summer internship by iVEC. The authors would like to acknowledge valuable discussions with Chris Harris at iVEC@UWA, and also thank the referees for their constructive comments and suggestions. T.L. is supported by the International Postgraduate Research Scholarship, Australian Postgraduate Award and the Bruce and Betty Green Postgraduate Research Top-Up Scholarship.

References

[1] Y. G. Chen, J. B. Wang, Qcompiler: Quantum compilation with the CSD method, Computer Physics Communications 184 (2013) 853–865.

[2] M. A. Nielsen, I. L. Chuang, Quantum Computation and Quantum Information, Cambridge University Press, Cambridge ; New York, 2011.

[3] P. Shor, Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer, SIAM Journal on Computing 26 (1997) 1484–1509.

[4] L. M. K. Vandersypen, M. Steffen, G. Breyta, C. S. Yannoni, M. H. Sherwood, I. L. Chuang, Experimental realization of shor's quantum factoring algorithm using nuclear magnetic resonance, Nature 414 (2001) 883–887.

[5] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, H. Weinfurter, Elementary gates for quantum computation, Physical Review A 52 (1995) 3457–3467.

[6] D. Deutsch, A. Barenco, A. Ekert, Universality in quantum computation, Proceedings: Mathematical and Physical Sciences 449 (1995) 669–677.

[7] G. Cybenko, Reducing quantum computations to elementary unitary operations, Computing in Science Engineering 3 (2001) 27–32.

[8] R. R. Tucci, A rudimentary quantum compiler (2cnd ed.), arXiv preprint quant-ph/9902062 (1999).

[9] M. Mttnen, J. J. Vartiainen, V. Bergholm, M. M. Salomaa, Quantum circuits for general multiqubit gates, Physical Review Letters 93 (2004) 130502.

[10] V. Bergholm, J. J. Vartiainen, M. Mttnen, M. M. Salomaa, Quantum circuits with uniformly controlled one-qubit gates, Physical Review A 71 (2005) 052330.

[11] F. S. Khan, M. Perkowski, Synthesis of multi-qudit hybrid and d-valued quantum logic circuits by decomposition, Theoretical Computer Science 367 (2006) 336–346.

[12] K. Manouchehri, J. B. Wang, Quantum random walks without walking, Physical Review A 80 (2009) 060304.

[13] A. De Vos, Y. Van Rentergem, Multiple-valued reversible logic circuits, Multiple-Valued Logic and Soft Computing 15 (2009) 489–505.

[14] M. B. Alexis De Vos, Reversible computation, quantum computation, and computer architectures in between., Multiple-Valued Logic and Soft Computing 18 (2012) 67–81.

[15] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, LAPACK Users' Guide, Third Edition, Society for Industrial and Applied Mathematics, 1999.

[16] J. Kempe, Quantum random walks: an introductory overview, Contemporary Physics 44 (2003) 307327.

[17] S. D. Berry, J. B. Wang, Two-particle quantum walks: Entanglement and graph isomorphism testing, Physical Review A 83 (2011).

[18] T. Loke, J. B. Wang, Efficient circuit implementation of quantum walks on non-degree-regular graphs, Physical Review A 86 (2012).

[19] R. Cleve, A. Ekert, C. Macchiavello, M. Mosca, Quantum algorithms revisited, Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences 454 (1998) 339–354.

[20] B. Douglas, J. Wang, Efficient quantum circuit implementation of quantum walks, Physical Review A 79 (2009).